

SAR: Stroke Authorship Recognition

Sara Shaheen, Alyn Rockwood and Bernard Ghanem

King Abdullah University of Science and Technology (KAUST), Thuwal, Saudi Arabia
sara.shaheen@kaust.edu.sa, alynrock@gmail.com, Bernard.Ghanem@kaust.edu.sa

Abstract

Are simple strokes unique to the artist or designer who renders them? If so, can this idea be used to identify authorship or to classify artistic drawings? Also, could training methods be devised to develop particular styles? To answer these questions, we propose the Stroke Authorship Recognition (SAR) approach, a novel method that distinguishes the authorship of 2D digitized drawings. SAR converts a drawing into a histogram of stroke attributes that is discriminative of authorship. We provide extensive classification experiments on a large variety of data sets, which validate SAR's ability to distinguish unique authorship of artists and designers. We also demonstrate the usefulness of SAR in several applications including the detection of fraudulent sketches, the training and monitoring of artists in learning a particular new style and the first quantitative way to measure the quality of automatic sketch synthesis tools.

Keywords: sketch, stroke, stroke segments, style, authorship recognition, fraud detection, sketch training

ACM CCS: Categories and Subject Descriptors: image processing, computer vision — shape recognition

1. Introduction

Research in voice and face recognition has developed methods to recognize and distinguish individuals based on pertinent visual and audio cues [TEBEH06]. Inspired by this concept, we investigate whether individuals (specifically artists) are distinguishable by the way they draw, and if so, the extent to which this uniqueness can identify their drawings, so as to help train them or detect sketch fraud. This work is the first to address the problem of authorship recognition from drawings. Our new method, called stroke authorship recognition (denoted as SAR) assumes that an artist's style can be recognized by the frequency distribution of certain mathematical descriptors, which are deduced from basic sketching strokes. A fundamental assumption is that artist style manifests itself mathematically in the artist's basic strokes. For example, we observe that Disney's Mickey Mouse tends to exhibit rounded strokes using many, nearly elliptic curves, whereas Looney Tunes' Daffy Duck consists of a combination of straighter and tightly curved strokes. By extracting enough strokes from a digitized drawing, a characteristic histogram is created. This histogram captures the inherent style of the overall drawing. As illustrated in Figure 1, SAR consists of five steps: (1) obtain sketches, (2) extract and segment strokes from each sketch, (3) calculate a number of mathematical attributes for each

stroke segment, (4) use these attributes to represent each sketch as a distribution of feature frequencies and (5) use these features to train a machine to recognize the author.

SAR is useful in a variety of applications such as sketch fraud detection (i.e. discriminating fraudulent from original sketches), where there is a need to examine fine-grained stroke-level features to discriminate sketches that *look* similar. This problem cannot be approached using existing shape matching techniques, which are unable to detect local and fine differences. Also, these methods are heavily dependent on the content of the sketch and not its style. Moreover, SAR can play a major role in artistic style training and reproduction, since it can be used to quantitatively assess how the sketching style of an artist in-training is becoming similar to a desired target style throughout the training process. For example, newly recruited artists at Disney are typically required to undergo a 6-month training procedure to familiarize themselves with the company's sketch techniques. This is done to ensure that new and existing characters can be created with the same Disney *look*. Clearly, there is a need for an automated technique to examine how an artist progresses during his/her sketch style training. Another area that can make use of SAR is handwriting verification and signature analysis, where the current state of the art uses features that are centric/specific to handwriting and the focus is on the uniqueness

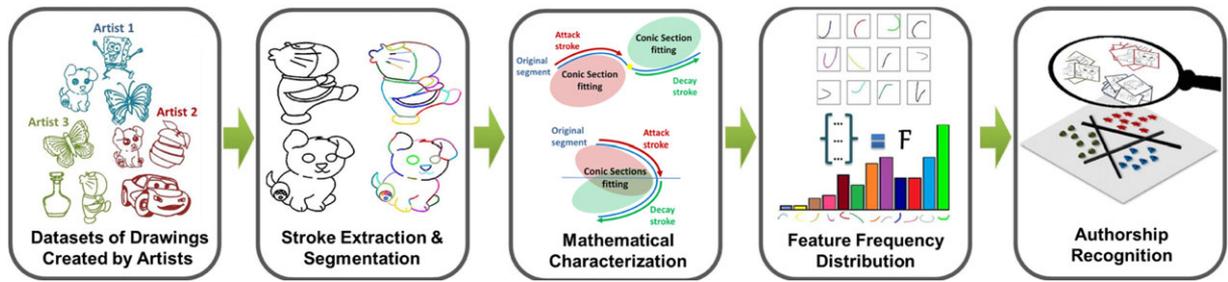


Figure 1: The SAR pipeline: (1) compile sketches from different artists; (2) extract strokes and split them into segments; (3) characterize each stroke segment mathematically; (4) represent each sketch as a distribution of feature frequency; (5) train a machine to recognize authorship based on (4).

of letters, punctuation formation, as well as, flow and structure. SAR, on the other hand, provides a more general context to investigate the uniqueness of drawn strokes. SAR is useful to other areas including design patent litigation and brand marking. As subtle differences exist among sketches of different artists, especially when attempting fraud, we focus on examining low-level stroke features to discriminate between styles that are very similar. These features are empirically validated and biologically inspired.

Contributions In this work, (1) we propose a novel method for sketch authorship recognition (SAR) based on the hypothesis that the collection of an artist's strokes in a sketch are unique to that artist and their frequency can be used to define his/her style. SAR can be used for multiple purposes including the detection of sketch fraud. (2) We compile two new sketch data sets collected from a number of experienced artists. These data sets are designed to expose SAR to different levels of challenges and sketch variations. They will be made publicly available, along with their digitized form and SAR source code, to allow for further research on this topic. (3) We develop two SAR-enabled sketch applications. The first provides artists in-training with immediate feedback on how close their sketching style has become to a particular target style and monitors their progress throughout the training process. The other provides the first quantitative and automatic measure to evaluate the quality of automatic sketch synthesis tools.

2. Related Work

In this section, we survey previous work that is most related to our approach, grouped into five main categories.

Shape Matching Determining stroke authorship seems akin to shape matching, but has different requirements as further inspection is needed. Numerous methods have been developed for shape matching and classification in the past [MM92, BMP01, JHY03, BBM05], as well as, recently [MOA11, IAP*11]. Shape matching searches for similar shapes between two images, where one is usually considered the query image. Our system differs from shape matching in two main aspects. First, shape matching focuses on global information of contours such as zero crossings of curvature, while SAR segments the contour and studies detailed local features from each curve segment, e.g. the eccentricity of a conic fit. Sec-

ond, shape matching is usually applied to low-resolution images for computational reasons, where these images tend to be classified by their content, e.g. a set of different mice silhouettes/shapes from different artists tend to belong to the same object class. Our SAR technique can be employed efficiently on images of various sizes, especially those at high resolution. Unlike shape matching, SAR is less dependent on sketch content and more focused on the local intricacies of sketch style.

Sketch and Artistic Style Analysis A data-driven approach in [BSM*13] is developed to analyse style in portrait sketches. This technique is also used for portrait sketch synthesis. While their focus is to mimic and synthesize a particular sketching style of only human portrait sketches, our focus is to discriminate sketches based on their authorship (i.e. artistic style). In their study, they analyse sketches at the level of strokes and shapes. However, they handle stroke analysis differently, since they focus on global features. SAR on the other hand exploits local stroke features that are necessary in discriminating authorship of sketches that look similar as is the case in sketch fraud detection. Moreover, in [BSM*13], portrait sketches were digitally collected using the Wacom pen to build a library of strokes while we allow artists to express their style freely and to erase or redraw their sketches using any medium they prefer. Interestingly, we use the sketch synthesis data set of [BSM*13] to demonstrate how SAR can quantitatively evaluate sketch synthesis. SAR results are on par with those reported by the authors after an online study with human participants.

In [LFTC13], a total of 13 000 drawings of faces was collected. Unlike our work, they do not study the problem of authorship as they focus on auto-correction of strokes for novice artists. In [LYFD12], the authors mimicked a particular artistic style via matching that was based on filtered velocities and shape context. Concurrently, the work of [KNBH12] synthesized new drawings using hatching styles of sketching that are learned by example, while that of [FTP03] provided an example-based method to modify line drawings for the purpose of reproducing different artistic styles. In [HOCS02], a method of learning statistical models of 2D curves is used to synthesize new curves based on an example style. Also, the approach of [CGL*08] studied where artists draw lines in sketches of objects such as tools, automobile parts and bones. They concluded that artists tend to draw similar lines in consistent locations. In our experiments, we use the data set of this work to show interesting

new results regarding the uniqueness of sketch style despite strict sketching constraints imposed on the artists. Finally, a large body of work focuses on generating artistically stylized rendering using 2D input images or videos of non-photorealistic rendering (NPR). We refer the reader to the survey of [KCWI13] for details. Although this work targets the analysis of sketches and artistic styles of different sketches, it does not address the important problem of how authorship can be determined based on stroke cues manifesting themselves in sketches.

As for sketch analysis, recent work [Gri11] provides suggestions for research directions in 3D sketching based on an observational study, where artists are asked to comment on how and why they made the marks on a number of objects they sketched. Also, a semi-automated method [NSS*13] was developed for semantic segmentation of digital sketch drawings to support further sketch analysis.

Sketch Recognition and Retrieval The work of [EHA12] developed an automated data-driven method to explore a large collection of hand-drawn sketches using drawings collected by many non-experts. Their primary goal was to represent sketch content to perform object recognition from a sketch and *not* sketch style. Concurrently, in [SWZZ12], a system was proposed for real-time recognition and retrieval of semantically meaningful attributes of hand-drawn sketches. Their work is not limited to pre-defined object classes. Other sketch retrieval methods are based only on geometric similarity between sketches [SMGE11, EHBA11]. Unlike our work, the field of sketch recognition and retrieval is based on classifying sketch content from a discrete number of semantic categories. They do not provide comparisons among similar sketches or across artistic styles and do not address the problem of authorship recognition.

Forensic Handwriting Analysis Forensic handwriting analysis is a well-studied problem [SL03] as is signature analysis, and more precisely, offline feature-based signature verification [IP08], [KC13]. Handwriting analysis tools tend to be very specific to the problem domain, so they use features centric to handwriting such as letter height, pitch, baselines, crossings, etc. They do not generalize well to sketch analysis [SS04]. Their focus is on the uniqueness of letter and punctuation formation, flow and structure [FK01], whereas we investigate the uniqueness of hand-drawn strokes in a more general context.

Graphical-Based User Authentication A number of graphical-based user authentication methods such as doodle sketches were proposed as an alternative to conventional authentication methods. As an example, PassDoodle is a graphical based authentication mechanism, which attempts to identify users by their handwritten designs (doodles). For authentication, the query doodle is represented on a regular grid and matched to training doodles in order to determine user authenticity [VKR04, GM07]. The effectiveness of using PassDoodles for user authentication is demonstrated here [Ren09]. In [OKX*08], on the other hand, a sketch-based authentication method was developed to find the *closest* sketch in a training database by extracting edge orientation features from a user's query sketch. Unlike SAR, these authentication methods neither build a classification model nor provide an analysis of different artistic styles. Moreover,

they do not build an intermediate sketch representation of the user input which makes them similar to shape matching techniques. An exposition of the shortcomings of graphical based-authentication methods is given in [Gan10].

In summary, our work provides novel analysis for different sketching styles using a representation of locally extracted features, which distinguish it from existing work. In this section, we presented an overview of the most relevant and recent work, all of which are areas where SAR can provide useful applications to. This ranges from shape matching to sketch analysis, recognition and retrieval as well as handwriting analysis and graphical-based user authentication.

3. Stroke Authorship Recognition

In this section, we give a detailed description of our proposed SAR approach. The overall pipeline of SAR is depicted in Figure 1. Given a sketch image, the first stage in SAR is to extract major stroke contours in the image and segment these strokes into stroke segments. These strokes occur both at the silhouette and the interior of the sketch. Stroke segments from many sketches across multiple artists are grouped (using low-level image features) to form a universal dictionary of stroke segments. This dictionary is employed in a hierarchical bag-of-words (BoW) model, which is used to represent a sketch image **I** as a histogram of stroke segments. We claim that this stroke histogram encodes some of the characteristics of an artist's unique style and thus can be used to discriminate this artist's sketches from others. Discrimination is performed in a supervised manner using multi-class classification, where each class designates an artist. In what follows, we elaborate on the individual stages of SAR.

3.1. Stroke extraction and segmentation

Image **I** is decomposed into strokes that are further split into stroke segments that can expose authorship. Representing sketches at such a local scale can identify stroke segments that play an essential role in discriminating authorship.

Stroke Extraction We use Adobe Live Trace, an off-the-shelf digitization technique, to decompose a sketch into a set of paths, each of which comprises a number of Bezier curves depicting artistic strokes [ADO10]. There exists a number of other commercial and non-commercial digitization techniques, such as the recent work by Noris *et al.* [NHS*13]. However, we choose Adobe Live Trace in SAR, since it stays faithful to the original sketch and it is widely accessible and easy to use. With Adobe Live Trace, we can apply different levels of digitization to the original sketch. Using this tool, we keep all the digitization parameters the same across all sketches of all artists, excluding the grey threshold value, which is varied among artists. The values of these parameters are set as follows: Blur = 0 px, Max Stroke Weight = 30 px, Min Stroke Length = 10 px, Path Fitting = 50 px, Corner Angle = 20° and noise = 2 px. This variation is needed because artists used different types of pens with different pressure. In Section 6.2, we evaluate the effect of this digitization on sketch style recognition. Extracted strokes are then traced and sampled as pixels in the image for stroke segmentation.

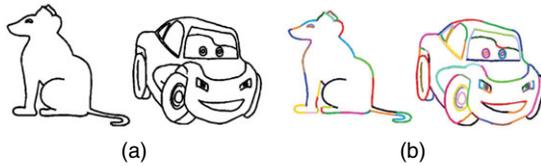


Figure 2: Stroke segmentation. (a) shows two sketches after digitization and (b) shows their extracted stroke segments (colour-coded).

Stroke Segmentation Strokes extracted from **I** can be quite long and contain a rich amount of geometric information. Modelling these strokes as a whole is a difficult task in itself, since it should encode variations that a particular stroke can take on. Instead, we resort to breaking each stroke into smaller units (called segments) that are represented in a more straightforward and conventional manner. We fit a b-spline curve to each stroke and identify breakpoints as locations of local maximum curvature in the b-spline. We explicitly handle linear strokes by placing breakpoints at its two ends to avoid oversegmentation. The pixels between two consecutive breakpoints (or the beginning/end of the stroke) are grouped together and denoted as a stroke segment. Figure 2 shows stroke segments extracted from two sketches taken from one of our sketch data sets.

3.2. Mathematical characterization

After **I** is decomposed into stroke segments, we represent it according to its stroke content. We aim to describe a stroke segment's structure and the manner in which it is drawn. We focus on local features that are simple to extract, representative and invariant to various deformations (e.g. rotation, translation and scale). In this paper, each segment is described by four simple features that encode eccentricity, symmetry, local consistency and inflection. The first three features describe the stroke segment itself, while the last one describe its relationship to its neighbours. These features are both empirically validated and biologically inspired. When considering the process needed to analyse drawing traits at the stroke level, one fundamental factor arises, namely, the neuro-geometry of hand-arm movement, i.e. the physiology in drawing strokes. In the foundational paper of [Mor81], it was shown that different subjects produce hand/arm movements that are very similar to or coincide with simple curve strokes. Although low curvature change was generally maintained across different subjects and experiments, it was also found that individuals exhibited unique stroke characteristics. This finding was also confirmed in the work of [ABM82]. In [FH85], a mathematical model based on minimizing jerk (change of curvature) and analysing stroke velocity was verified empirically. In summary, this and other work (see references in [FH85]) strongly suggest that (1) gesturing (i.e. the physical mechanism of drawing) and strokes can be unique and identifiable for an individual, and (2) that curvature and its change are key features for analysing simple strokes. Inspired by the above physiological findings, we characterize stroke segments by focusing on curvature and its change in each stroke segment.

Moreover, in analysing a large number of simple strokes, we see that a conic fit to a stroke segment is faithful to its original geometry and representative enough for the purpose of authorship recogni-

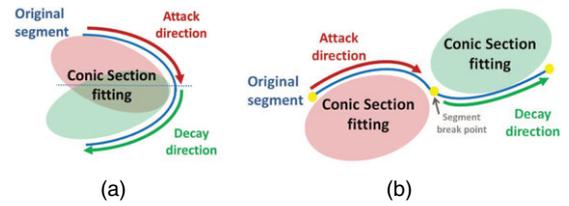


Figure 3: Attack and decay parts used to construct stroke segments. Part (a) shows a stroke segment divided into equal length parts enabling the computation of the symmetry feature. Part (b) shows two adjacent stroke segments connected by an inflection breakpoint. The inflection feature is computed by comparing these two segments.

tion. From each conic, we extract eccentricity, symmetry local consistency, and inflection features that are conveniently invariant to rigid-body deformations. Needless to say, this invariance also holds when the entire sketch is represented using stroke segment features.

Eccentricity: This feature represents the change of curvature in a stroke segment. It measures how the segment deviates from being circular, i.e. how rounded or sharp it is. By fitting a conic section to the stroke segment [Tau91], eccentricity (ϵ) is computed as a function of the ratio between the conic's major and minor axis lengths as in (1).

$$\epsilon = \begin{cases} \sqrt{1 - \frac{b^2}{a^2}} & \text{if the conic is an ellipse} \\ \sqrt{1 + \frac{b^2}{a^2}} & \text{if the conic is a hyperbola.} \end{cases} \quad (1)$$

Symmetry: To crudely evaluate how symmetric a stroke segment is, we simply take the absolute difference in eccentricity of the segment's two halves (refer to Figure 3a). The segment is divided into two equal length parts (denoted as *attack* and *decay* following typical sketch nomenclature) and the eccentricity of each part is computed as described above.

Local Consistency: This feature measures the extent of local variation within a stroke segment. It is computed as the average absolute difference between the eccentricity of the entire segment and the eccentricity of distinct overlapping parts of this segment. To generate these distinct parts, we employ a sliding window approach across the segment, where the window size is one-third the size of the segment itself and the step size is half the window size. This feature captures subtle changes in shape and curvature.

Inflection: The previous three features describe the stroke segment itself. However, characteristics of artistic style are encoded in how stroke segments are sequenced. Of special interest are locations of inflection, where the sign of curvature changes. An inflection point is detected by SAR as a breakpoint between two stroke segments that together form a stroke (refer to Figure 3b for an example). To encode such relational information at an inflection point, we compute the absolute difference in eccentricity between each pair of stroke segments that share this inflection point. For segments forming a T-junction (i.e. a stroke segment with inflection points at either end), we compute their average eccentricity difference. For

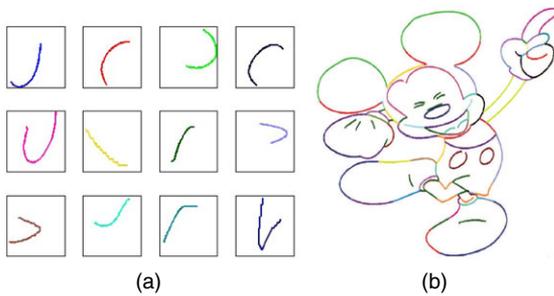


Figure 4: (a) Some elements of the universal stroke segment dictionary constructed from the segments of the training sketches. (b) Colour-coded nearest neighbour assignment of segments to dictionary elements.

stroke segments whose breakpoints are not inflection points, we set their inflection feature to a nominal value (-1).

3.3. Feature frequency distribution

Having each stroke segment s_i of sketch image \mathbf{I} characterized by the 4D feature vector described above, this stage represents \mathbf{I} as a distribution of frequencies of those features.

Building a Universal Stroke Segment Dictionary Our assumption is that the frequency in which an artist uses particular types of strokes is a suitable indicator of his/her authorship. To formalize this observation, we compile a dictionary of stroke segments that tends to be universal among different artists. The dictionary is learned on the stroke segments of the sketches that are used in training only. To construct a dictionary of n elements, we apply hierarchical k -means clustering on a large set of stroke segments. In our experiments, we set $n = 60$. We denote this dictionary as the universal stroke segment dictionary as it tends to capture the most commonly used stroke segments among artists. This clustering step is the first stage of the BoW model popularly used to represent natural images for image classification [SZ03]. We use hierarchical k -means, since it is more robust and less sensitive to the choice of n than traditional k -means clustering. And it generates a tree of n cluster centres (and not only a set of centres as in the case of k -means), which encodes the membership of any stroke segment at all levels of the tree and not just the leaves. To encode a stroke segment, one has to *traverse* this tree starting from the root and recursively select among its children the nearest cluster centre in 4D feature space. As a result, each stroke segment is encoded as a binary membership vector of length n , where each value reflects whether the feature traversed the corresponding tree node or not. In Figure 4, we show image examples of cluster centres in the universal stroke segment dictionary, as well as, the colour-coded membership of each stroke segment in a sketch. Such assignment is independent of where a stroke segment exists in a sketch.

BoW Features After constructing the universal dictionary, we represent a sketch as a histogram of stroke segments following the traditional BoW framework. Sketch image \mathbf{I} containing m stroke segments produces m binary membership vectors each of length n . These membership vectors are pooled together to produce the BoW

feature vector \mathbf{f}_I . We use Gaussian weighted mean pooling [Bou14] to encode the frequency in which an artist uses each element of the universal dictionary. As we will see, coupling \mathbf{f}_I with a discriminative model enables authorship recognition. Since the original four-stroke segment features are invariant to rotation, translation and scale, the BoW feature vector is invariant to these deformations [LW07].

3.4. Authorship recognition

So far, a sketch image is represented by a sparse n -dimensional histogram depicting the frequency in which each type of stroke segment is used by the artist in the sketch. As such, we expect the BoW feature to possess enough discriminative power to determine authorship based on simple strokes alone. We validate this assumption empirically in Section 6.1. Given a training set of sketches labelled according to the artist who drew them, we build a discriminative model using the training BoW features. In order to reduce testing time and to pinpoint the most discriminative portions of the BoW feature, we employ a forward feature selection procedure, which greedily appends a single feature at a time *only* if this addition improves classification accuracy [JZ97]. Based on our experiments, only a small subset of the n features is actually used for discrimination. Using these subsampled BoW histograms, we build a multi-class classifier to assign authorship to an unseen sketch image. For sketch fraud detection, we use a binary classifier (original vs. fraudulent) instead. We experimented with a variety of classifiers and found that an RBF (Radial Basis Function) kernel SVM and a kNN classifier achieve the best performance. However, we decided on the kNN classifier because of its minimal training time. We use two experimental setups to evaluate the accuracy of our classifier: two-fold cross-validation and leave-one-out. The first is a popular setup in image classification, while the latter sheds light on how dependent the classifier performance is to the amount of training data. Following common practise, the best k value for the kNN classifier is estimated using cross-validation on a finite set of possible k values.

4. Sketch Data Sets

To evaluate the performance of SAR, we compile two new sketch data sets collected from experienced artists. Participating artists were screened, so as to guarantee high levels of sketching capability. The chosen artists were graphical or interior designers by profession, each with 7–10 years of sketching experience. We gave them the freedom to draw using a pen, pencil or digitally at any scale. Moreover, they were allowed to correct their strokes or redraw the entire sketch with no time limit. Since we had direct access to these professionals, we were able to administer different sketching scenarios. This allowed us to control the *content variety* (i.e. what the artist draws in a sketch) and the extent of *sketching constraints* (i.e. how the artist should sketch). These two factors impact the strokes an artist chooses and ultimately his/her style. To our knowledge, this work is the first to compile such a diverse set of sketch data for the purpose of studying authorship from strokes. To allow for further research in this area, we will make all these data sets (images and annotations) publicly available. Next, we describe the data sets used in this paper (refer to Table 1 for a summary).

Table 1: A summary of data sets used in this paper. We compiled the first two and reused the line study data set of [CGL*08].

Data set	No. Artists	No. Sketches	No. Strokes	Content variety	Sketching constraints level
Free Style	10	100	5500	Ten different sketching objects	No constraints (Draw this sketch)
Fraud	9	90	4900	Ten different sketching objects	High constraints (Draw and simulate fraud)
Line Study	11	107	6000	Up to 12 different objects	Very high constraints (Redraw over the original image)

Free Style Data Set We collected 10 images of objects from diverse semantic categories and then asked 10 artists to sketch them. The objects were chosen to be detailed enough to adequately reflect artistic style and to be relatively easy to replicate by an experienced artist in a reasonable amount of time. On average, it took each artist 2 weeks to submit the 10 sketched objects. In total, the 10 artists collectively drew 100 sketches. No constraints or specific instructions were given to the artists, thus, giving them complete sketching freedom.

Fraud Data Set Here, we simulate a sketch fraud scenario. Using the same images in the free style data set, we chose the sketches of one of the artists as *original* drawings. We asked the other nine artists to draw all the original sketches. We provided them with an instruction sheet, where they were requested to draw the original sketches in such a way that it would be very hard to distinguish their *own* drawings from the originals, thus, simulating sketch fraud. The nine artists were prohibited from using methods or supplies (e.g. translucent tracing paper) that could produce copies of the original sketches. Each artist took an average of 2 weeks to finish the copying task. We collected a total of 90 new sketches. (refer to Figure 5 for examples). Since the artists were requested to simulate sketch fraud, this data set is highly constrained from an artistic style point of view. Therefore, it poses a substantial challenge to any authorship recognition system, albeit manual or automatic.

Line Study Data Set We also use the line study data set generated in [CGL*08]. This data set is a collection of sketches of various objects (e.g. mechanical tools, automobile parts and bones) drawn by multiple artists. The sketching tasks were highly constrained, since artists were asked to draw their sketched lines over a faded copy of the original image. Clearly, the task of recognizing sketch authorship in this data set is more difficult than the previous two. We selected all sketches from artists, who drew at least six sketches, thus, leading to a data set of 107 sketches from 11 artists.

5. Human Sketch Authorship Recognition

For comparison and as a baseline, we study the inherent difficulty of the authorship recognition problem for humans. The human visual system (HVS) is renowned for its effectiveness in successfully completing many high-level recognition tasks (e.g. object and action recognition), so much so, that it remains the *gold standard* to which automated recognition systems aspire. Despite significant advances in computer vision, the HVS almost always outperforms

automated methods in such tasks. However, there *do exist* recognition tasks, in which the HVS underperforms. These tasks usually deal with *fine-grained* recognition of objects (e.g. faces), where the interclass variation is minimal and on par with the intraclass variation. We motivate this fact with an example. Although the HVS can easily discriminate between a ‘chair’ and a ‘dog’, it does not do so well in recognizing a person’s face in a large data set of people (e.g. the entire population of a country) from the same cultural and racial background. The differences between people’s faces are so subtle that the minute details discriminating them are hard to uncover by the HVS. However, it is exactly these details that automated recognition methods focus on. This enables them to outperform the HVS in these tasks (e.g. robust face recognition [WYG*09]). In this section, we provide extensive empirical evidence from two user studies that highlight the sheer difficulty people and experienced artists face when trying to recognize authorship from 2D sketches.

Authorship Recognition by Non-Artists The aim of this study is to shed light on how accurately people can recognize sketch authorship. As an online quiz, participants are first shown sample sketches from the free style data set labelled with artists who drew them. We showed five randomly selected sketches from each of the 10 artists. Next, we administered for each online participant a total of five questions, each of which asked him/her to assign an *not previously seen* sketch to one of the 10 artists that he/she thought drew it. We chose to show users only 5 of the 10 images per artist so as not to overwhelm them and to enable a fair comparison with the two-fold experiments we conducted using our automated method. We allowed participants to zoom-in to the images when needed. To reach a large number of participants, we published our user study on Amazon Mechanical Turk (AMT). To validate the quality and seriousness of each Turker’s answers (as usually done in practise), we administer control questions at the beginning of the quiz. Moreover, Turkers are randomly directed to one of the many versions of the online study and answers from unique workers are recorded.

After 6 weeks of activity and more than 2000 unique participants, the accumulated results of this study show that people can successfully recognize the authorship of a sketch (among 10 different artists) with an average accuracy of 29% and with standard deviation of 8%. Participants took an average of 4 min to complete the assigned quiz. We see that average human performance is moderately better than a random choice classifier (i.e. choose one of the 10 artists at random irrespective of what the sketch is) with an average accuracy of 10%. Obviously, this establishes that sketch

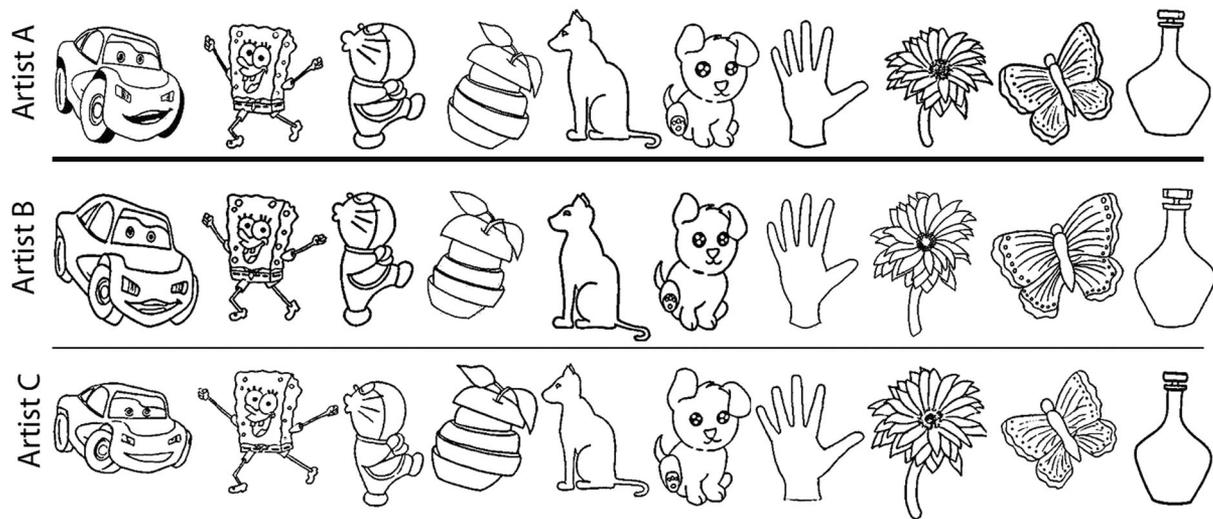


Figure 5: Samples of sketches from the fraud data set. The first row shows the original sketches drawn by Artist A. The other two rows show the sketches of two other artists (Artists B and C), who attempted sketch fraud. Obviously, the fraudulent sketches look extremely similar to the originals, making authorship recognition (manual or automatic) quite difficult.

authorship recognition is quite difficult for people. More importantly, we show later that our automated SAR method achieves an accuracy of 56% under the same conditions.

Sketch Fraud Detection by Artists We conduct a similar study, but we target artists specifically because it is conceivable that the average person might find this task difficult due to his/her lack of sketching experience and/or artistic talent. The aim here is to evaluate the performance of artists in discriminating fraudulent sketches from original ones. A total of 25 experienced artists were given an online quiz, where five original and five fraudulent sketches were made known to each user. The sketches were taken from the fraud data set described earlier. Each artist is then given a set of five recognition questions, whereby he/she needs to label the unseen sketch as original or fraudulent. Moreover, artists were given the chance to view their quiz results and to e-mail us feedback.

The quiz results show that the artists could distinguish original sketches from fraudulent ones only 47% of the time, as compared to 50% random chance. Each artist took an average of 5 min to complete the task. As expected, most of the feedback we received elaborated on how truly difficult the task is. Later, we show that SAR significantly outperforms the artists' recognition accuracy, thus, motivating the plausibility of using SAR in an automated fraud detection system for sketches (e.g. patent drawings and cartoon sketches). We include links to both quizzes in the Supporting Information.

6. Experimental Results and Evaluation

This section presents a number of experimental results we obtained to assess multiple facets of the proposed SAR approach. (1) To evaluate the effect of sketching constraints on SAR, as well as, its overall effectiveness in recognizing sketch authorship and detecting sketch fraud, we test SAR on the datasets described in Section 4. (2) We also evaluate the sensitivity of SAR to a number of algorithmic variations.

These results, in turn, are used to design the most representative and discriminative variant of SAR.

6.1. Computational authorship recognition

We ask the following questions: Are simple strokes unique to the artist who draws them? If they are, then to what extent can they identify the author who drew them? To our knowledge, these questions are not adequately addressed in prior work. To answer them, we conduct several experiments on three separate data sets, which provide a rich diversity of sketch content and sketching constraints. In what follows, we report SAR classification accuracy for each data set and provide cross-data set analysis. We also evaluate SAR's ability to detect sketch fraud using the fraud data set.

Free Style Data Set On this constraint-free data set, SAR accuracy on test data is 56% and 53% using leave-one-out and two-fold validation, respectively, with 7% standard deviation. For comparison, random chance in this case is 10%. To verify the statistical significance of our classification results, we run a standard student *t*-test with a significance level of 95%. The null hypothesis, 'SAR performs similar to a random classifier', is rejected 97% of the time (i.e. *p*-value is 0.03).

Fraud Data Set. Is artistic style among different artists still preserved when they are consciously attempting to commit fraud? To answer this, we use our fraud data set to build a 10-class SAR classifier to discriminate among the 10 artists. Using this data set, random chance is 10%. Here, SAR reaches a significantly higher average accuracy of 48% (leave-one-out) and 43% (two-fold cross-validation) with 6% standard deviation. This result suggests that an artist's sketching style is still distinguishable (to a certain extent) even when attempting sketching fraud. To visualize the discrimination power of SAR, we show the BoW features of three sketches (same object) drawn by three artists from this data set in Figure 6.

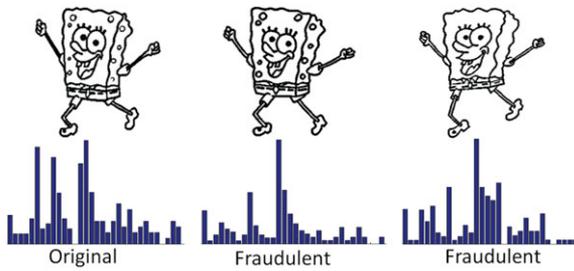


Figure 6: BoW features of the same sketch object drawn by three artists (one original and two fraudulent). Although they look very similar, their features are different.

Table 2: A summary of SAR performance on three data sets. Average accuracy (%) is reported for leave-one-out and two-fold setups. Random chance accuracy (%) is also reported for comparison.

	Rand chance	Leave-one-out	Two-fold	std
Free style	10	56	53	7
Fraud	10	48	43	6
Line study	9	33	30	7

One of the sketches is an original and the rest are fraudulent. Although all three sketches look very similar, their underlying features have obvious differences, which SAR capitalizes on to determine authorship. By conducting the student t -test with a 95% significance level, we conclude that SAR is more discriminative than a random classifier with a p -value of 0.04.

Line Study Data Set The sketching task in the data set compiled by [CGL*08] is highly constrained as discussed in Section 4. As it is similar in spirit to the fraud data set, we expect similar results. We train and test SAR in the same way as before but with a total of 11 artists and 107 sketches. With 11 classes, random chance is 9%, which is significantly lower than SAR's average accuracy of 33% (leave-one-out) and 30% (two-fold) with 5% standard deviation. This noteworthy discrepancy in performance indicates that artists still maintain some uniqueness in their sketching style, even though they are forced to copy a different style altogether. The p -value of the student t -test in this case is 0.05.

We give a performance summary for SAR on the three data sets in Table 2. The slight difference in accuracy between leave-one-out (i.e. all but one sketch are used for training) and two-fold cross-validation (i.e. only 50% of sketches are used for training) indicates that SAR is reasonably insensitive to the amount of training data used. This also suggests that SAR is generalizable, a promising property for a classifier in the presence of unseen data.

Cross-Data set Analysis In the data sets above, the contributing artists were subject to varying levels of sketching constraints, ranging from unconstrained (free style data set) to highly constrained (line study data set). To investigate the sensitivity of SAR performance with sketching constraints, we compare its accuracy across

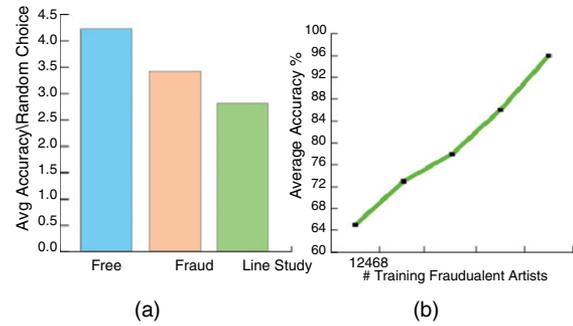


Figure 7: (a) The ratio between SAR accuracy and random chance on three data sets organized in increasing order of sketching constraints. Clearly, the more constraints imposed on artists, the less discriminative their strokes become. (b) Fraud recognition performance improves as more fraudulent artists are used in training, especially on sketches from artists who were not included in training.

all data sets in Figure 7(a). The data sets are ordered in increasing levels of constraint. To normalize the effect of different numbers of classes across data sets, we plot the ratio of SAR accuracy to random chance. As expected, recognizing sketch authorship becomes harder as more constraints are imposed on the artist. However, on all the data sets, SAR performance is significantly (at least 2.5 times) higher than random chance. This suggests that artists do *not* lose all characteristics of their unique style even under the strictest of constraints.

Fraud Recognition Experiment Here, we setup a fraud recognition experiment (original vs. fraudulent) that predicts SAR performance in a real-world scenario. It is unrealistic to assume that the SAR classifier will have access to fraudulent training examples from *all* artists. However, we expect that when more fraudulent samples are used in training, SAR's test performance on fraudulent sketches from artists, who were *not* included in training, will improve. In other words, knowing more about how fraud *looks* like will help SAR better detect sketch fraud, even for fraudulent artists whose sketches are not trained on. To validate this expectation, we train the SAR binary (fraud vs. original) classifier with fraudulent sketches from increasingly more fraudulent artists (i.e. from one to five artists). In each case, 50% of the original sketches are used for training (two-fold validation). Then, each SAR classifier is tested on the remaining original and fraudulent sketches. This means that SAR will always be tested on fraudulent sketches from artists, whose style has not been seen during training. In Figure 7(b), we plot the average test accuracy of SAR as the number of fraudulent artists used in training increases. Clearly, this result endorses our aforementioned expectation. In fact, SAR's fraud detection performance reaches 96% when eight of the nine fraudulent artists are used for training and the ninth for testing. Despite the extremely high similarity between fraudulent and original sketches, we conclude that SAR is able to effectively spot artistic fraud. Moreover, this result suggests that SAR can be used in an online active learning setup, where test samples are sequentially classified and in turn weighted and then used in retraining the classifier.

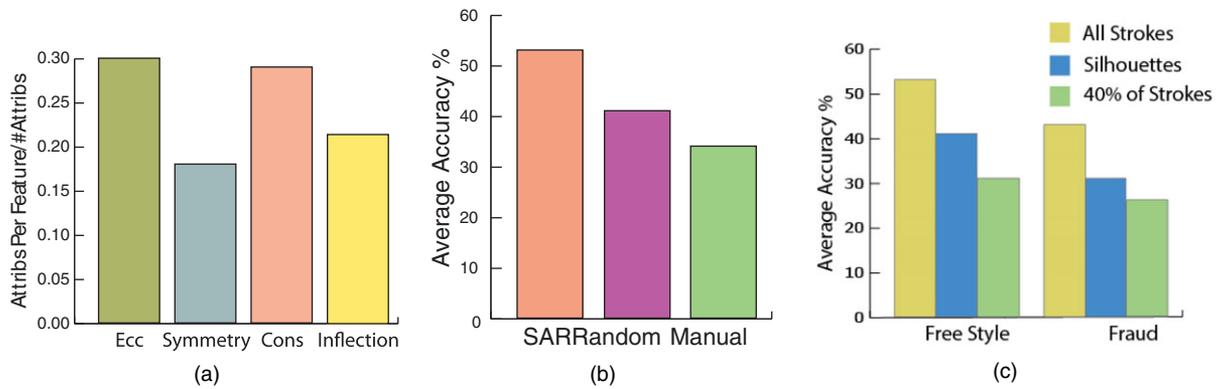


Figure 8: (a) Contribution of each feature type to SAR recognition. (b) SAR accuracy on the free style data set when its segmentation method is replaced with a manual or random one. SAR accuracy drops by 15% when the segmentation is done manually and 20% when it is done randomly. (c) SAR accuracy on the free style data set, when the stroke segments are taken to be all the segments in the entire sketch, all the segments in the silhouette of the sketch or a randomly sampled fixed percentage of the segments in the entire sketch. In the last case, we only consider the percentage of segments in the sketch that are from the silhouette, which is estimated to be around 40% for each sketch. Adding internal stroke segments does improve accuracy but the improvement is only about 10%. Silhouette strokes are more discriminative than the same number strokes sampled from the entire sketch. This shows the discriminative power of the silhouette.

6.2. Algorithmic details

Varying the details of the different computational modules of SAR (refer to Figure 1), leads to a variant of SAR. In this section, we investigate many of these variations and provide algorithmic details to reproduce the SAR classifier.

Feature Contribution Stroke segments are represented by four biologically inspired features, as discussed in Section 3.3. Here, we study the contribution of each feature to SAR’s discriminative power. To do this, we employ a conventional forward feature selection method that greedily determines which features (along with their importance weights) should be incorporated into the SAR classifier. This is a data-driven process, so different training-test splits of the data set can lead to different feature selections and weights. By training SAR (with two-fold validation) multiple times on the free style data set, we accumulate an average selection weight for each feature (refer to Figure 8a). We see that eccentricity and local consistency are the most discriminative followed by inflection and symmetry with no one feature dominating.

Variations in Stroke Segmentation To justify our stroke segmentation method (refer to Section 3.1), we compare it against two baseline methods: manual and random stroke segmentation. For manual segmentation, we asked 10 artists to manually segment sketches into strokes as they see fit, while in the random case, breakpoints are selected uniformly at random within each extracted stroke. We set a minimum length for each stroke segment to prevent singularities. SAR accuracy (two-fold) using each of the three segmentation methods on the free style data set is summarized in Figure 8(b). Clearly, our automated segmentation method outperforms the other two. Similar to the human recognition result in Section 5, SAR improvement over manual segmentation lends more evidence to the fact that human performance is suboptimal in fine-grained tasks, such as authorship recognition.

Silhouettes versus All Strokes In many cases (e.g. the Mickey Mouse cartoon character), the silhouette of a sketch can be very discriminative, so much so, that silhouettes have been used as primary features in other recognition tasks (e.g. action recognition [LZL08]). Here, we study how discriminative silhouette stroke segments are on their own within the SAR pipeline. In Figure 8(c), we use two data sets to compare SAR accuracy when the stroke segments are taken to be all the silhouette segments, all the sketch segments or a sample set of the sketch segments, where the sample size is based on the percentage of silhouette stroke segments in the sketch. This percentage is 40% on average. Interestingly, silhouettes seem to be quite discriminative in their own right, as SAR performance remains high even when only silhouettes are used. In fact, including stroke segments from the sketch interior only improves accuracy by about 10%. When comparing silhouette segments to the same number of segments sampled from the entire sketch, we see that the former is much more discriminative, thus, validating the discriminative power of a sketch’s silhouette.

Effects of Digitization As mentioned in Section 3.1, stroke segments are extracted from a sketch using the well-known Adobe Live Trace digitization technique. The level of digitization is controlled by the user to trade-off faithfulness to the original sketch with compactness of representation. For instance, high levels of digitization tend to smooth out strokes, thus, masking aspects of the artist’s intrinsic style and possibly affecting SAR performance. To investigate the effect of digitization on SAR, it is applied to the fraud data set after applying three distinct levels of digitization. This is done by varying the path fitting parameter in Live Trace. We report the average accuracies due to this parameter variation in Table 3. As expected, the higher the digitization level, the lower the SAR accuracy is. This result suggests another application that could benefit from SAR, namely, the quantitative assessment of various digitization techniques. SAR can conceivably be used to evaluate how much of the unique sketching style is preserved after digitization.

Table 3: Effects of digitization on SAR accuracy. Sketches from the fraud data set are processed using three levels of digitization in Adobe Live Trace. SAR accuracy decreases significantly with higher levels of digitization.

Level of digitization	Leave-one-out	Two-fold cross-validation
Low	48	43
Medium	40	36
High	28	23

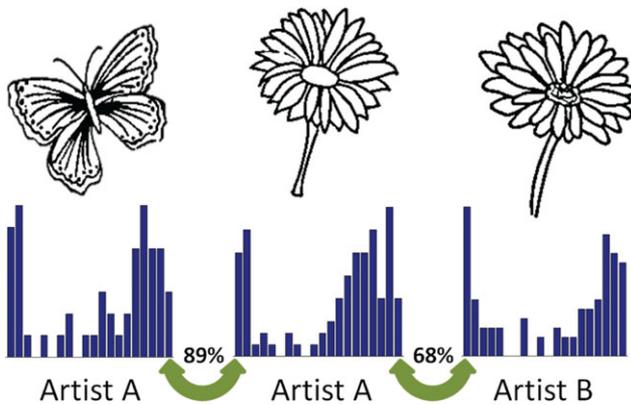


Figure 9: A visual comparison between the BoW features of three sketches from the free-style data set: a flower drawn by Artist A (middle), a butterfly drawn by Artist A (left) and the same flower object drawn by Artist B (right). From the features themselves and their average pairwise similarity scores, we see that SAR discrimination abstracts artistic style and is not significantly affected by sketch content.

Style not Content SAR is designed to represent and classify artistic sketch style, while not being significantly sensitive to sketch content in general. This is a primary factor that distinguishes SAR from existing shape matching techniques discussed in Section 2. In fact, our experiments show that sketches of different content (objects) drawn by the same artist tend to have more similar BoW features than sketches of the same content drawn by different artists. For instance, the BoW feature of a flower drawn by one artist is more similar to a butterfly drawn by that artist than the feature of the same flower drawn by another artist, as illustrated in Figure 9. To quantify this observation, we use a normalized Gaussian similarity measure $s(\mathbf{I}_1, \mathbf{I}_2)$ based on the Euclidean distance between two BoW features. From the free style data set, we randomly select a sketch \mathbf{I}_i^A drawn by artist A and the same sketch \mathbf{I}_i^B drawn by another artist B to compute $s(\mathbf{I}_i^A, \mathbf{I}_i^B)$. We then randomly select a sketch \mathbf{I}_j^A drawn by A with $i \neq j$ and compute $s(\mathbf{I}_i^A, \mathbf{I}_j^A)$. By performing this operation across the whole data set, we compute the average interartist similarity $s(\mathbf{I}_i^A, \mathbf{I}_i^B) = 68\%$ and the average intra-artist similarity $s(\mathbf{I}_i^A, \mathbf{I}_j^A) = 89\%$. This result shows that SAR can abstract an artist's style from different sketch content. We would like, however, to emphasize that strokes are not completely independent from the content. For example, SAR cannot recognize the authorship of sketches of a single square if an author's train-

ing data *only* includes circles. Therefore, SAR inherently assumes that the training data contain non-trivial sketches with sufficient stroke diversity characterizing the overall style of any particular artist.

6.3. Human versus SAR

In this section, we summarize the difference in recognition performance between human subjects and our SAR automated system.

we compare human performance against our authorship recognition computational model. The first user study was built using the free style sketching data set and our user study shows that people can successfully recognize the authorship of a sketch among seven different artists with an average accuracy of 36% and with a standard deviation of 10%, using a sample of more than 2000 participants. SAR, on the hand, recorded an accuracy of 60% and a standard deviation of 8% when applied to the same data set and using the same setup as in the user study. We also conducted another user study which involved 25 artists and used the fraud experiment database. Artists were asked to distinguish fraudulent and original sketches using samples of both. User study results showed that given a set of sample sketches, artists can distinguish between fraudulent and original sketches with an average accuracy of 52% and a standard deviation of 10% while SAR classified fraudulent and original sketches under a setup similar to the user study with 96% accuracy and with a standard deviation of 14%. As can be seen, SAR outperformed human performance in classifying the authorship of similar sketches of which we conclude that SAR is predicted to be a useful tool to assist people and professionals in distinguishing authorship of sketches.

7. Applications

In this section, we show SAR's feasibility in two useful applications other than sketch fraud detection, namely, sketch style training and sketch synthesis evaluation.

7.1. SAR for sketch style training

How persistent are artistic styles when artists undergo extensive training for the explicit purpose of altering their style and adopting another? For example, this type of training is administered in major cartoon companies (e.g. Walt Disney). For this purpose, we develop a SAR-based application that allows artists, designers and animators in-training to quantify their progress in adopting the *target* style of a particular artist, whose sketches have been encoded in SAR as demonstrated next.

Style Training Setup We first determined a target style by finding a reasonably known artist, who has published his work online along with specific instructions on how to follow his drawing style. The target artist also provided YouTube videos showing how his sketches can be drawn step-by-step. We selected five sketches from his portfolio, specifically those with video instructions. Then, we identified three artists with varying levels of sketching and artistic experience (novice, intermediate and advanced) to undergo the style training process. The advanced level artist is a professional with 10 years of



Figure 10: Samples sketches drawn by artists at the first and final stages of training along with the target style sketches. Sketches on the left correspond to the intermediate artist while ones on the right to the novice artist. The progress in their sketching styles is noticeable for both artists but more so for the novice one.

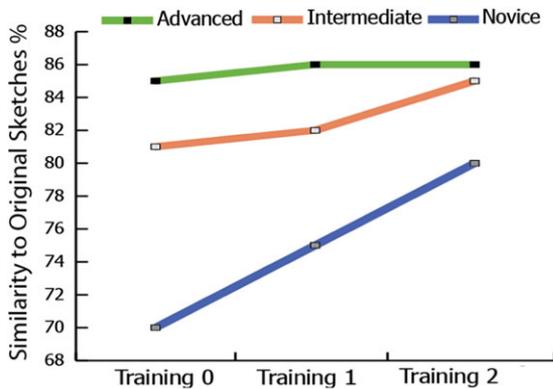


Figure 11: SAR similarity-to-target scores for 3 novice, intermediate, and advanced level artists after 3 stages of training. The novice artist has the least score in the initial stage but progresses the most with training.

experience and the intermediate artist has 3, while the novice artist only sketches as a hobby.

In the first stage of training, we asked all three artists to draw the five target sketches and then used the SAR-based tool to give them quantitative feedback on how similar their sketching style is to the target. The measure based on BoW features defined earlier in Section 6.2 is used to compute style similarity. In the second stage, we asked the artists to draw the sketches again after consulting a set of textual instructions on how to draw the target sketches. These instructions were obtained from the target artist's website. In the third stage, we provided the artists with their new SAR similarity scores along with instructional online videos. After submitting their newest sketches, we again computed their SAR similarity scores. Sample sketches of both novice and intermediate-level artists at the first and last stages of training are shown in Figure 10. It is obvious from these results that the novice artist has improved more significantly than the intermediate artist in adopting the target style.

Style Training Results Figure 11 plots the evolution of the SAR similarity score throughout the three-stage training process for each artist. As expected, each trainee's score increases with training. The novice artist exhibits an overall improvement of 10%, which can be attributed to the explicit use of textual and video instruction for

target training. The initial similarity scores for the intermediate and advanced-level artists are much higher than the novice one, but they also exhibit an improvement of 4% and 2%, respectively. The slight improvement by the advanced artist conveys how difficult it is for an artist with a well-defined style and extensive experience to adopt a new style, as well as, how much easier it is for a novice artist to do the same. Based on our results, we conclude that SAR can be successfully employed to quantitatively monitor training progress. A training executable is available in the Supporting Information. We will also release the training source code to extend this application to a larger set of sketches from more participating artists and to allow users to build their own training applications using their own sketch data.

7.2. SAR for evaluating sketch synthesis results

As mentioned in Section 2, there is a number of recent methods that focus on sketch synthesis and artistic style analysis. Their aim is to develop automatic sketch synthesis tools that mimic a particular artistic style. Most of these tools do not provide a quantitative assessment of how close the synthesized style is to the target one, thus, making it difficult to compare synthesis methods. A few of them have validated their synthesis quality through extensive user studies, which are cumbersome to compile and require careful analysis. We propose using SAR to quantitatively evaluate sketch synthesis methods without the need to conduct tedious user studies. To exemplify this application, we use SAR to analyse the synthesis results of [BSM*13]. In their work, portrait sketches are synthesized from artistic styles of seven artists, each of which drew 24 portrait sketches. The real and synthesized sketches are publicly available.

We first use SAR to classify artistic style among the real portrait sketches only. In this case, SAR accuracy is 52% (leave-one-out) and 50% (two-fold), as compared to random chance accuracy of 14%. When classifying only synthesized sketches, we obtain 54% accuracy (leave-one-out) and 51% accuracy (two-fold), respectively. The similarity in performance between the two scenarios suggests that the synthesized portraits are as hard to distinguish as the real ones and that the synthesis tool of [BSM*13] maintains a very similar amount of style diversity as in the real sketches. Next, we use SAR to classify the real sketches using the synthesized ones only as training and vice versa. In this case, SAR accuracy dropped to 26% (leave-one-out) and 25% (two-fold). This performance *drop* can be used as a quantitative measure of synthesis quality, as it allows for comparison between different synthesis methods. In fact, an ideal

synthesis tool should produce a SAR accuracy drop close to zero, since the style of the synthesized sketches should be indistinguishable from that of the real sketches.

Our results are on par with those of the three user (perceptual) studies in [BSM*13]. In fact, we setup the SAR classification experiments above to follow the same setup used in these user studies, where the only difference between the two is that SAR performs the classification automatically without any human feedback. The participants in the user studies registered a very similar accuracy when classifying the real and synthesized portraits separately (as did SAR). More importantly, the drop in classification by human participants is around 20%, which is comparable to the 25% drop reported by SAR. This result validates that SAR can be used to automatically and quantitatively evaluate sketch synthesis tools, without the need to conduct cumbersome user studies.

8. Conclusions and Future Work

In this paper, we shed light on a new direction in sketch analysis, namely, authorship recognition through stroke analysis. We propose a SAR approach that discriminates between artistic sketch styles based on the choice and frequency of use of basic strokes. SAR is motivated by real-world issues, including the fact that fraudulent copies of original sketches with patent and copyright protection are sold on some websites, which claim them to be original. As it is hard for an individual to discriminate authentic and fraudulent sketches (Section 5), original sketches from known companies and artists are in danger of replication and copyright infringement. Moreover, cartoon artists undergo months of training before they can officially contribute, so as to ensure that the style in which the characters are drawn is not distorted when artists are replaced (Section 7.1). Similar to handwriting and signatures, sketches can be used for user authentication as discussed in Section 2. Thus, detecting major and minor differences in sketching style is important. Furthermore, design patents that contain sketches of commercial designs need to be protected, as is apparent in the ongoing lawsuits between Apple and Samsung for example. All of this shows the need to have automatic tools, such as SAR, to assess the authenticity of sketched objects and identify their authorship.

From our extensive experiments and user studies (Section 6.1), we provide empirical evidence that SAR *does* encode unique and consistent characteristics of an artist's sketching style, which are in turn used to discriminate one artist's sketches from others. We also showed that the extent to which SAR can be used for discrimination is dependent on the sketching constraints imposed on the artist. Moreover, our extensive user studies empirically validate the difficulty of this fine-grained recognition problem and indicate that SAR can improve upon human performance. All the compiled data (data sets and user studies) and source code will be made publicly available to enable further research on this exciting topic and to allow for quantitative comparison.

Future Work We aim to improve the discriminative nature of SAR by enhancing the learned universal stroke segment dictionary. One way to do this is to investigate how strokes are actually generated by the artists through tracking their hand movements while sketching. We believe this will provide us with more information about

an artist's sketching style. Although the segmentation process is currently viewed as a pre-processing step in SAR, we aim to investigate how supervised information (artist labels) can be incorporated in this process so as to produce stroke segments that are inherently discriminative. We also aim to consider the way artists draw by collecting sketches using a digitizing pen so we can take advantage of the pressure, tilt and speed used during the creation of the drawing.

Acknowledgement

This work was supported by competitive research funding from King Abdullah University of Science and Technology (KAUST).

References

- [ABM82] ABEND W., BIZZI E., MORASSO P.: Human arm trajectory formation. *Brain* (1982).
- [ADO10] ADOBE: Illustrator.
- [BBM05] BERG A., BERG T., MALIK J.: Shape matching and object recognition using low distortion correspondences. In *Proceedings of CVPR* (2005), vol. 1, IEEE, pp. 26–33.
- [BMP01] BELONGIE S., MALIK J., PUZICHA J.: Matching shapes. In *Proceedings of ICCV* (2001), IEEE, vol. 1, pp. 454–461.
- [Bou14] BOUREAU Y.-L.: *Learning Hierarchical Feature Extractors for Image Recognition*. PhD thesis, New York University, 2014.
- [BSM*13] BERGER I., SHAMIR A., MAHLER M., CARTER E., HODGINS J.: Style and abstraction in portrait sketching. *ACM Transactions on Graphics* 32 (2013), 55:1–55:12.
- [CGL*08] COLE F., GOLOVINSKIY A., LIMPAECHER A., BARROS H. S., FINKELSTEIN A., FUNKHOUSER T., RUSINKIEWICZ S.: Where do people draw lines? *ACM Transactions on Graphics* 27 (2008), 88:1–88:11.
- [EHA12] EITZ M., HAYS J., ALEXA M.: How do humans sketch objects? *ACM Transactions on Graphics* (2012).
- [EHBA11] EITZ M., HILDEBRAND K., BOUBEKEUR T., ALEXA M.: Sketch-based image retrieval: Benchmark and bag-of-features descriptors. *IEEE Transactions on Visualization and Computer Graphics* 17 (2011), 1624–1636.
- [FH85] FLASH T., HOGAN N.: The coordination of arm movements: An experimentally confirmed mathematical model. *Journal of Neuroscience* 5 (1985), 1688–1703.
- [FK01] FRANKE K., KPPEN M.: A computer-based system to support forensic studies on handwritten documents. *International Journal on Document Analysis and Recognition* 3, 4 (2001), 218–231.
- [FTP03] FREEMAN W. T., TENENBAUM J. B., PASZTOR E. C.: Learning style translation for the lines of a drawing. *ACM Transactions on Graphics* (2003), 33–46.

- [Gan10] GANI A.: A new algorithm on graphical user authentication (GUA) based on multi-line grids. *Scientific Research and Essays* (2010), 3865–3875.
- [GM07] GOVINDARAJULU N. S., MADHVANATH S.: Password management using doodles. In *ICMI* (2007).
- [Gri11] GRIMM C.: *Results of an observational study on sketching*. Tech. rep. WUCSE-2011-57, Washington University in St. Louis, 2011.
- [HOCS02] HERTZMANN A., OLIVER N., CURLESS B., SEITZ S. M.: Curve analogies. In *Proceedings of Rendering Techniques* (2002), pp. 233–246.
- [IAP*11] ION A., ARTNER N. M., PEYRÉ G., KROPATSCH W. G., COHEN L. D.: Matching 2D and 3D articulated shapes using the eccentricity transform. *Computer Vision and Image Understanding* (2011).
- [IP08] IMPEDOVO D., PIRLO G.: Automatic signature verification: The state of the art. *IEEE Transactions on Systems, Man, and Cybernetics* (2008), 817–834.
- [JHY03] JIN R., HAUPTMANN A., YAN R.: Image classification using a bigram model. In *Proceedings of AAAI Spring Symposium on Intelligent Multimedia Knowledge Management* (2003).
- [JZ97] JAIN A., ZONGKER D.: Feature selection: Evaluation, application, and small sample performance. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 19, 2 (1997), 153–158.
- [KC13] KOVARI B., CHARAF H.: A study on the consistency and significance of local features in off-line signature verification. *Pattern Recognition* 34, 3 (2013), 247–255.
- [KCWI13] KYPRIANIDIS J. E., COLLOMOSSE J., WANG T., ISENBERG T.: State of the ‘art’: A taxonomy of artistic stylization techniques for images and video. *IEEE Transactions on Visualization and Computer Graphics* 19, 5 (2013), 866–885.
- [KNBH12] KALOGERAKIS E., NOWROUZEZAHRAI D., BRESLAV S., HERTZMANN A.: Learning hatching for pen-and-ink illustration of surfaces. *ACM Transactions on Graphics* 31 (2012), 1:1–1:17.
- [LFTC13] LIMPAECHER A., FELTMAN N., TREUILLE A., COHEN M.: Real-time drawing assistance through crowdsourcing. *ACM Transactions on Graphics* 32, 4 (2013), 54:1–54:8.
- [LW07] LU D., WENG Q.: A survey of image classification methods and techniques for improving classification performance. *International Journal of Remote Sensing* 28, 5 (2007), 823–870.
- [LYFD12] LU J., YU F., FINKELSTEIN A., DiVERDI S.: Helping-hand: Example-based stroke stylization. In *ACM Transactions on Graphics* (2012), vol. 31, pp. 46:1–46:10.
- [LZL08] LI W., ZHANG Z., LIU Z.: Expandable data-driven graphical modeling of human actions based on salient postures. *IEEE Transactions on Circuits and Systems for Video Technology* 18, 11 (2008), 1499–1510.
- [MM92] MOKHTARIAN F., MACKWORTH A.: A theory of multiscale, curvature-based shape representation for planar curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14 (1992), 789–805.
- [MOA11] MICHEL D., OIKONOMIDIS I., ARGYROS A.: Scale invariant and deformation tolerant partial shape matching. *Image and Vision Computing* 29 (2011), 459–469.
- [Mor81] MORASSO P.: Spatial control of arm movements. *Experimental Brain Research* 29 (1981), 459–469.
- [NHS*13] NORIS G., HORNUNG A., SUMNER R. W., SIMMONS M., GROSS M.: Topology-driven vectorization of clean line drawings. *ACM Transactions on Graphics* (2013), 4.
- [NSS*13] NORRIS G., SYKORA D., SHAMIR A., COROS S., HORNUNG A., SUMNER R., SIMMONS M., WHITED B., GROSS M.: Smart scribbles for sketch segmentation, Sept. 19, 2013. US Patent App. 13/424,083.
- [OKX*08] OKA M., KATO K., XU Y., LIANG L., WEN F.: Scribble-a-secret: Similarity-based password authentication using sketches. In *Proceedings of ICPR* (2008), IEEE, pp. 1–4.
- [Ren09] RENAUD K.: On user involvement in production of images used in visual authentication. *Journal of Visual Languages and Computing* (2009), 1–15.
- [SL03] SRIHARI S. N., LEEDHAM G.: A survey of computer methods in forensic document examination. pp. 278–281
- [SMGE11] SHRIVASTAVA A., MALISIEWICZ T., GUPTA A., EFROS A. A.: Data-driven visual similarity for cross-domain image matching. In *Proceedings of ACM SIGGRAPH* (2011).
- [SS04] SRIHARI S., SHI Z.: Forensic handwritten document retrieval system. In *Proceedings of International Workshop on Document Image Analysis for Libraries* (2004), IEEE, pp. 188–194.
- [SWZZ12] SUN Z., WANG C., ZHANG L., ZHANG L.: Sketch2tag: Automatic hand-drawn sketch recognition. In *Proceedings of ACM MM* (2012).
- [SZ03] SIVIC J., ZISSERMAN A.: Video Google: A text retrieval approach to object matching in videos. In *Proceedings of ICCV* (2003).
- [Tau91] TAUBIN G.: Estimation of planar curves, surfaces, and non-planar space curves defined by implicit equations with applications to edge and range image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 13 (1991), 1115–1138.
- [TEBEH06] TOLBA A., EL-BAZ A., EL-HARBY A.: Face recognition: A literature review. *International Journal of Signal Processing* 2, 2 (2006), 88–103.

[VKR04] VARENHORST C., KLEEK M., RUDOLPH L.: Passdoodles: A lightweight authentication method. Research Science Institute, 2004.

[WYG*09] WRIGHT J., YANG A., GANESH A., SASTRY S., MA Y.: Robust face recognition via sparse representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 31 (2009), 210–227.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's web site:

<http://ivul.kaust.edu.sa/Documents/Data/SAR-Authorship-Recognition-CGF.zip>